

Compliments of **IBM**

IBM Limited Edition

Flash Array Deployment

FOR
DUMMIES[®]
A Wiley Brand

Learn:

- Why data storage performance matters to your business
- How flash storage benefits you
- Why all-flash arrays offer more value
- Which IBM FlashSystem deployment design is best for you

Neal Ekker



***Flash Array
Deployment***
FOR
DUMMIES®
A Wiley Brand

IBM Limited Edition

By Neal Ekker

FOR
DUMMIES®
A Wiley Brand

Flash Array Deployment For Dummies®, IBM Limited Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2015 by John Wiley & Sons, Inc.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. IBM, the IBM logo, FlashSystem, FlashCore technology, Real-time Compression, and Variable Stripe RAID are trademarks or registered trademarks of International Business Machines Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub.

For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-119-10219-9 (pbk); ISBN: 978-1-119-10208-3 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Carrie A. Johnson

Editorial Manager: Rev Mengle

Acquisitions Editor: Steve Hayes

Business Development Representative:
Sue Blessing

Key IBM contributors: Matthew S. Key,
Philip A. Clark, Arati V. Vora

Table of Contents

Introduction	1
About This Book	1
Icons Used in This Book	2
Chapter 1: Learning Why Data Storage Performance Matters	3
Speed Matters Online	3
eCommerce	5
Big Data and analytics	6
Financial services	7
Cloud, mobile, and social engagement	7
Performance Drives Value	8
Chapter 2: Getting to Know Flash Storage Systems	11
Solid State Drives	12
SANs and Storage Arrays	14
PCIe Cards	16
Solid State Arrays	17
Chapter 3: Choosing Flash Storage Arrays	21
Storage Analysis	21
Flash Controllers	23
Looking at the Challenges of SSDs	25
Understanding the Liabilities of PCIe Cards	26
Establishing the Advantages of Flash Arrays	27
IBM FlashSystem	28
Chapter 4: Exploring Deployment Designs with IBM FlashSystem	31
Deploying Direct Attached Storage with IBM FlashSystem	32
Connecting Systems with a SAN-Attached Architecture	35
Chapter 5: Implementing the Future with Virtualized Storage	37
Storage Pooling and Tiering	38
Data Protection	40
Capacity Optimization	43

Introduction

Organizations ranging from city governments, through movie production companies and medical research institutions, to car manufacturers are realizing that how well their information technology (IT) performs strongly affects how well their business performs. Insightful executives and IT managers understand that their data storage systems play a crucial role in how well their information technology helps them achieve important objectives such as faster decision making, better customer service, or a smaller data center budget. Solid state storage made from NAND flash memory chips has evolved in terms of cost, performance, and reliability to the point where many organizations are seriously considering its use to replace inefficient, unacceptably slow mechanical spinning disk systems. This accelerating trend has led enterprises to ask some natural questions: When should flash be used? Which flash solution is best for each particular use case? And how can I make it a successful, cost-effective part of my data center? These are the questions I answer (especially the last one) in *Flash Array Deployment For Dummies*, IBM Limited Edition.

About This Book

If you're a decision maker in an enterprise determined to make more, spend less, and move faster, this book is for you. *Flash Array Deployment For Dummies*, IBM Limited Edition, tackles the data storage challenges of "enterprises" — commercial, scientific, and governmental organizations. It does *not* address "consumer" data storage issues, such as those faced by privately owned PC, smartphone, laptop, and iPad devices, and so on.

Chapters 1 and 2 of this book are most helpful to decision makers. In these chapters, I introduce some of the data storage-related problems you see that have led you to consider flash storage, discuss why you may choose to solve these problems with flash storage, and highlight some benefits if you do.

Also in the early chapters, I introduce the various types of flash storage and explain what they're used for and who currently uses them. As the name implies, this book is ultimately about *flash storage arrays* — flash devices that can stand alone and are used most often in data center environments where multiple computers (servers) can access or share the same storage solution. Flash arrays offer good solutions to the majority of storage challenges you may be experiencing in your data center.

Chapters 3 and 4 provide the most current thinking about what you should do as the responsible manager or technician if you are assigned the task of actually implementing a flash storage solution. Of course, this information can be invaluable to those working on the data center floor. But it may also prove helpful to IT decision makers because how effectively your flash storage solution is deployed, configured, and operated will play a large role in the return you see in your flash storage investment — a matter dear to the hearts, and careers, of IT decision makers.

Icons Used in This Book

You'll find several icons in the margins of this book. Here's what they mean.



A Tip is a suggestion or a recommendation. It usually points out a quick and easy way to get things done or provides a handy piece of extra information.



The Warning icon alerts you to conditions that require extra care and thinking. For example, you don't want to omit critical steps in evaluating your needs and planning your implementation.



Anything that has a Remember icon is something that you want to keep in mind.



Technical Stuff contains information that's interesting and useful but not vital to understanding flash array deployment. Info here may include a brief history of a principle, the earliest practitioners, or the origin of a word. It also showcases technical points. You can either read these or skip over them.

Chapter 1

Learning Why Data Storage Performance Matters

.....

In This Chapter

- ▶ Learning why storage speed matters to online applications
 - ▶ Seeing how storage speed accelerates traditional IT environments
-

Information technology (IT) isn't an end in itself; it's a means to solving certain business problems or enhancing business opportunities. A data storage solution that makes life better in the data center but doesn't contribute positively to your organization's success isn't really a solution at all. So, the first order of business in learning more about flash array deployments is to connect your data storage to your business challenges.

Speed Matters Online

Let's say that you're an enterprise that accomplishes at least some of your business activities online. Many organizations fall into this group — everyone from retailers selling products directly to customers through the Internet, to banks and other financial institutions offering services online, to scientific organizations sharing research information with colleagues around the world. Yet the Internet, at its most basic level, is just a collection of computers exchanging digital bits — ones and zeroes, pulses of high and low electrical voltage. Computers connected to each other and exchanging information are known as networks. Such collections or networks of computers have grown to be very complex and powerful over the years in terms of how much digital information they can

transfer over various connective media, such as metal wires, optical fibers, and even various frequencies of electromagnetic waves (wireless), and how well they can manage the streams of information zooming between them. They need to be powerful because the amount of data that you may want to transfer across various digital networks is rapidly increasing.



Around 2.5 quintillion bytes (a unit of digital information most commonly consisting of eight bits) of new data is created *every day* — and by 2017, IBM predicts that data volumes will grow by another 800 percent.

In order for computer networks to manage and transfer this ever-growing enormous amount of data successfully, they must become ever faster. Data is constantly stored on and retrieved from many of the computers or “nodes” involved in networks, and the speed of data storage directly affects the overall performance of things you want to do using computer networks.

Simple transfers of information from one computer to another are just the tip of the iceberg of reasons why people use networks. They often access computer software over the Internet and other types of networks, such as Local Area Networks (LAN) within or controlled by organizations or geographically wider LANs called Wide Area Networks (WAN). Network-connected applications allow you to share photos with your loved ones, transfer money from one account to another, and even hold business meetings with colleagues around the world. But none of this happens effectively without fast data storage and retrieval.

Essentially all business, government, and research activities in the modern world use computer applications as foundational tools. As the volume of data increases, the applications on which you depend must grow ever faster. In order for applications to perform more work for you in shorter time frames, they must store or write data to storage devices and retrieve or read this data in the least amount of time possible. The amount of time taken for data storage round trips, essentially the storage response times, is known as *storage latency*.



One of the most important limits or throttles on an application's ability to perform useful work is the storage latency of the computer system on which the application runs or is hosted. Storage latency is a central and crucial concept within this book.

There is another kind of latency within digital systems — network latency. When applications operate partly or mostly over networks, network latency is added to the computational and storage latencies that exist where the application is actually running, so network-based applications must address both their own local latencies and the network latency as well. Dealing with the local storage latency is bad enough; when you add network latency, the challenges multiply.

The applications employed as vital tools by your enterprise are limited in their performance by the latency and other performance-related characteristics of the associated computer system's data storage devices and designs or architectures. Addressing this one issue itself has spawned large and thriving industries within the world of information technology. Then, if your organization's activities utilize computer networks — the Internet or your own LANs or WAN — the issues associated with application performance, storage latency, and the need for speed grow even more thorny and challenging.

But where there is challenge, there may also be opportunity. Whole new industries and sectors of economic endeavor have been created by the advent of networked computing, and most others have been transformed or at least significantly affected by these technologies. In almost every case, fast data storage is a fundamental requirement for success.

eCommerce

Online retail activities, or what is called *eCommerce*, has become a driving force of global economics. And eCommerce provides an excellent example of why data storage performance matters. Business to customer sales facilitated by the Internet surpassed \$1 trillion several years ago, according to IBM research, and will soon account for over 5 percent of all worldwide economic activity. In addition to creating entirely new business models, eCommerce also competes directly with traditional brick and mortar stores. For example, currently in the United States 70 percent of consumers experience their first interaction with a brand online, and within a few years 50 percent of all retail dollars spent in the United States will be partly or entirely transacted digitally.

Online shoppers don't want to wait for information to make their buying decisions; they demand rapid responses. Rich

and complex web pages with dynamic content take longer to load, especially on mobile devices, and that creates a fundamental challenge for eCommerce providers.

The influence of eCommerce is growing in the overall marketplace by about 20 percent every year. As utilization and market opportunity escalate, the IT backbone supporting eCommerce becomes a critical path component to success. Retail websites that deliver information about products and interact with customers quickly and reliably usually increase both market share and profitability.

Big Data and analytics

Enterprises today collect tremendous volumes of data that's generated by a wide range of sources often at extreme velocities. These massive data sets are called *Big Data*. Discovering and communicating meaningful patterns in these large collections of data is called *Big Data Analytics*.

For all businesses, data itself is one of your most valuable assets, and Big Data Analytics may already be one of the most powerful new tools you use to gain competitive advantage, increase sales, and protect your business from fraud. But, the near real-time analysis and response velocities of Big Data Analytics require a storage environment with the lowest possible system latency. And the rapid transfer of enormous data volumes requires extraordinary storage bandwidth. Therefore, storage performance truly matters in any enterprise hoping to harness the benefits of Big Data.

Science needs performance too

Non-profit enterprises also mine Big Data for value. Take the Large Hadron Collider (LHC) for example, outside Cern, Switzerland. LHC experiments involve about 150 million sensors delivering data 40 million times per second. The data flow could exceed 150 million petabytes

annually, or around 500 quintillion (5×10^{20}) bytes per day — almost 200 times higher than all other data sources in the world combined. From analyzing that data, LHC scientists found glimmers of evidence for the existence of the Higgs Boson or “God particle.”

Financial services

The financial services industry offers another good example of why storage performance matters. Few industries have been so affected and accelerated by the Internet as the financial services sector, especially equities trading. Core banking systems are getting faster as they turn from pure systems of record to systems of customer engagement, with new online and mobile access rates increasing dramatically. On the securities side of the house, equities sales occur in milliseconds now days. Throughout the financial sector these trends lead to fierce competition where the performance of IT infrastructure makes the difference between the firms that capture market share and profits and those that don't. System latency and scalability are of critical importance to applications in this environment. Beyond operational transaction processing, risk and market assessment requirements of financial services enterprises have also fostered the industry-wide adoption of online analytical processing (OLAP) tools, further fueling the requirement for very fast IT systems and high-performance data storage.

Cloud, mobile, and social engagement

The future of online enterprise is a wild new world. The Internet has not only transformed traditional businesses, but also it's fostered the creation of entirely new industries, inspired new avenues for theft and crime, and even enabled a new model for delivering compute services themselves. This exploding new world of commerce and interaction, legal and otherwise, drives an arms race of new data storage technologies and solutions, all based on the ever-accelerating need for speed.

Mobile computing and online social engagement are two of these entirely new enterprises spawned by the Internet. Literally, they're already profoundly changing the arena of global business and society. Proliferating mobile technology and the spread of social business are empowering people with knowledge, enriching them through networks, and changing their expectations. For example, 57 percent of companies now expect to devote more than a quarter of their IT spending to mobile and social systems of engagement by 2016, nearly twice the levels of 2013.

At the same time that industries and professions are being remade by the Internet, the IT infrastructure of the world is being transformed by the emergence of Cloud computing. In all Cloud delivery models the IT infrastructure challenges related to data storage are similar. Most importantly, because applications and functionality delivered through a Cloud model come to end-users through networks, local or Internet or both, system latency is a critical issue. Overall response time includes both network latency and response delays generated at the compute source. Networks are growing ever faster, which shifts much of the focus on reducing latency to the data center itself and from there directly to the storage systems. This is why flash array deployments in Cloud and other network-centric environments are escalating, because only high-performance data storage can enable the future.



IBM estimates that by 2016, more than one-fourth of the world's applications will be available in the Cloud, and 85 percent of new software is now being built for Cloud compute environments. The delivery of IT as online services is creating new business models that are generating a market expected to reach \$250 billion in 2015.

Performance Drives Value

Not all business, governmental, and scientific activity happens online. If 5 percent of worldwide economic activity is facilitated by the Internet, then 95 percent is not. This suggests there's a lot of data processing that isn't network-enabled and instead happening locally within the physical walls of enterprises.

Does this "locally" occurring computer activity need fast data storage? And do the organizations that depend on computer programs as crucial tools in their operations really benefit much from revved up IT infrastructure? In fact, it's easy to show that both are very true.

Databases offer the most widely applicable example in the "local processing" category. Databases and database management systems (DBMS) have been around since the dawn of the information age.



A *database* is an organized collection of data typically used to model aspects of reality in a way that supports processes that require information, for example, modelling the availability of rooms in hotels in a way that supports finding a hotel with vacancies. DBMSs are computer software applications that interact with users, other applications, and the database itself to capture and analyze data. A general-purpose DBMS is designed to allow the definition, creation, querying, update, and administration of databases.

For example, if at work you use applications for financial accounting, to manage employee records, to provide customer service, or to track parts inventories or product supplies, you use applications that rely heavily on databases. Because databases are involved with the majority of applications, and data processing is involved in the majority of economic activity on the planet, it's not a stretch to suggest that the performance of databases affects and influences the activities of business, government, and science more than almost any other information technology.



Do databases benefit from storage that's faster than traditional spinning disks? A recent study conducted by analysts at Wikibon suggests that overall IT costs can be dramatically lowered by replacing conventional disk-based storage with much higher performance storage:

- ✓ 54 percent lower overall IT infrastructure cost
- ✓ 94 percent less administration and operational support outlays
- ✓ 76 percent reduction in environmental (power/space) expenses
- ✓ 52 percent lower software costs

But how did we leap from exploring how databases can perform better with faster storage to lowering the costs associated with the entire data center? It turns out that the performance of your data storage system dramatically impacts the costs of deriving value from your information in a number of ways:

- ✓ When you do more work for the same cost, the expense per unit of work goes down. Faster storage, with lower latency, enables databases to respond more swiftly to each request for data from users or applications.

- ✔ Over the decades, most IT components have relentlessly grown faster and more powerful, but traditional spinning hard disk drives have not, so the speed of data processors, often called CPUs (central processing unit), is now much greater than disk-based storage — on the order of 10,000 times faster (GHz/ms). This means that processors hosting applications may spin idly through many cycles waiting for a data request as it travels from CPU to database to storage and back again. In fact, utilization rates can hover below 10 percent in many traditional data centers. But add faster storage, and CPU utilization rates can shoot up much higher, depending on how well other system components are optimized. Higher CPU utilization rates create efficiencies all through IT systems. Fewer servers can be used to accomplish the same amount of work. And less software may be needed, running on those servers.
- ✔ It turns out that faster storage isn't mechanical storage, and with no moving parts, only electrons, it consumes much less electricity. This also means it throws off proportionately less heat, which translates into less air cooling needed in the data center. Over the past few years, environmental costs such as power and HVAC, and even the value of data center floor space, have risen dramatically. They now figure prominently into any accurate assessment of enterprise data storage costs.
- ✔ Mechanical components, such as spinning disk drives, tend to break down or wear out faster than electronic circuitry. This leads to many more repair and reconfigure episodes for database and system administrators. Plus, the systems designed to mitigate the slow performance of traditional storage often grow very complex and demand plenty of attention. It's easy to see that a very wide range of enterprises — from retail websites, through banks and stock traders, to staff managing personnel records and product inventories in businesses of all sizes — can benefit from and in fact are demanding faster data storage performance. And it's also clear that faster storage isn't based on spinning mechanical disks.

Even with the advent of the Internet, mobile apps, social engagement, and Cloud computing, databases are still involved in the vast majority of data processing. It's easy to see that the benefits they derive from high-performance storage will drive flash array deployments for years into the future.

Chapter 2

Getting to Know Flash Storage Systems

.....

In This Chapter

- ▶ Understanding solid state drives
 - ▶ Discovering storage arrays and storage area networks
 - ▶ Using PCIe cards
 - ▶ Defining solid state arrays
-

Chapter 1 demonstrated that faster data storage for enterprises of all types and sizes is a matter of cost, productivity, and competitive advantage. After you make the decision to implement faster data storage for your enterprise, you're ready to move to the second step in the process of flash array deployment, which is learning about what higher performance data storage technologies are currently available to you. And that's what I cover in *this* chapter.

Essentially, there's only one viable option right now — data storage made from integrated circuits instead of spinning disks. For decades it has been known as *solid state storage*. Originally, solid state storage consisted of random access memory (RAM) chips aggregated into large integrated groups or arrays. These devices were used as massive temporary holding places for data, called *buffers* or *caches*. Because RAM loses data when the power goes off, these devices relied on other components such as redundant power supplies and even batteries within their deployment environments to prevent data loss. Various solid state storage devices made from DRAM (dynamic RAM, the current version of RAM) are still available, though they're very expensive. They're used in

unique situations where the lowest possible storage latency is demanded, no matter the cost.

In the past ten years, solid state storage made from a type of integrated circuitry called flash memory has become very popular in consumer electronics because a chip the size of your thumbnail can hold quite a lot of data, even when you turn off the device or the battery goes dead. The use of flash memory chips in a wide and ever-growing spectrum of consumer products has driven their cost steadily lower over the past decade while spurring plenty of innovative engineering. As the cost has fallen and the capabilities and endurance have dramatically risen, flash has become viable for the more demanding IT environments found in modern enterprises.



Introduced by Toshiba in 1984, flash memory cells are made of “floating gate” transistors. NAND flash memory chips are composed of literally millions of flash cells and form the basis of devices built for storing the data generated by business, government, academic, medical, and scientific enterprises of all types from around the globe. When you compare the operational expenses (electricity, cooling, floor space), management costs, server and software outlays, and performance value of conventional disk systems and flash storage, you see that even though until recently the purchase price or dollars per gigabyte (\$/GB) of flash devices for enterprise use has been considerably higher, all the other costs can be quite a bit lower, making the overall costs much more equivalent. And now that the \$/GB are converging as well, the total cost of ownership (TCO) of the two storage types is tipping toward flash. Because of this, IT industry analysts predict that deployments of flash data storage solutions for enterprise use cases will dramatically rise in the next few years.

Solid State Drives

In the early 1990s, following the invention of flash memory, a new kind of solid state storage product evolved. Because disk-based systems were the most widely used enterprise storage solutions, only storage products that could operate in the hard disk drive bays of servers or the disk enclosures of enterprise storage arrays were practical. And the solid state drive (SSD) was born.

For nearly 20 years, the term *solid state disk* was used by most industry insiders to refer to any solid state storage device, no matter its shape (form factor), what it was made from, or how it was used. In the past few years, as the industry has grown and products have proliferated, the term *solid state disk* has fallen out of use and now SSD refers specifically to solid state storage products with hard disk drive form factors that interface with storage systems by using industry standard hard disk drive software or protocols. Such a device can be seen in Figure 2-1.



Figure 2-1: An example of an SSD.

SSDs are products of convenience, cost, and trade-offs. Because they're designed to connect to storage systems just like traditional disk drives do, they offer a convenient way for enterprises to add some solid state storage performance to conventional disk-dominated environments. As the gap between the speed and performance of CPUs and disk storage grew steadily wider over the years, enterprises had stronger motivation to look for something that could practically integrate with their existing storage systems but provide more inputs/outputs per second (IOPS) and lower latency than hard disk drives. SSDs filled the bill.



Additionally, SSDs offer another powerful advantage — the cost of a unit is considerably lower than other solid state storage devices. It's important to note that this doesn't necessarily mean that the cost per usable storage capacity or the cost per application transaction capability is lower for SSDs. A useful analogy might be that of a Chevy pickup versus a semi-truck and trailer. It's quite possible that because the semi can haul a load of many tons whereas the pickup can barely haul one ton, the cost per unit of hauling capacity could actually be lower with the semi. But buying a semi costs much more than buying a Chevy pickup.

So, in all the applications where a Chevy will do just fine, SSDs have flourished. On the consumer side, SSDs will fit in your PC or your laptop without your having to buy added equipment or software to handle them, for the most part. The same is true with both enterprise servers and conventional storage systems. In fact, in the past ten years most enterprise operating systems, virtualization software, and storage array controllers have been upgraded to handle SSDs, with varying degrees of effectiveness. Thanks to their convenience and cost per unit, SSDs have maintained their rank over the years as the hottest selling solid state storage devices.

SANs and Storage Arrays

An advantage of SSDs is their ease of deployment in conventional enterprise storage environments — both in servers and in the large collections of hard disk drives known as enterprise storage arrays that are deployed in Storage Area Networks (SAN). A common SAN is shown in Figure 2-2.

A SAN is the standard way in today's data centers to share a storage resource such as a storage array with multiple servers. A SAN is created by networking the servers by way of Fibre Channel or other connectivity through switches to one or more storage devices. Each storage device could be a large enterprise storage array with several, dozens, or even scores of individual SSDs involved.




Figure 2-2: A common SAN.



An enterprise storage array is a group of integrated hard disk drives or other storage media devices uses a computer known as a *controller* to manage collective activities. Figure 2-3 shows the disk enclosures, controllers, network switches, and other related hardware of an enterprise storage array all housed in a single cabinet.



Figure 2-3: An enterprise storage array with all its components in a single cabinet.

Your SAN could be composed of multiple storage arrays and other storage devices, such as tape drives and even optical/CD drives often used for data archiving purposes, as well as flash arrays. All these storage devices are connected using an appropriate networking technology such as Fibre Channel or Ethernet and then made available to your various application hosts/servers on the other side of a network switch.



SANs have proven to be powerful and popular storage designs, or architectures, for decades, but they do introduce network latency. Every request for data made by an application must travel away from the CPU, out of the server enclosure, through Fibre Channel or other networking, through a switch(s), into a storage array controller, round and about

through the storage array software and hardware, into the individual SSD where the data actually resides — then make the entire return journey.

PCIe Cards

By the early 2000s, enough applications were becoming storage latency sensitive that finding some alternatives began to look like an excellent business venture. Engineers explored ways to avoid the network latency incurred by traditional SANs, and soon the Peripheral Component Interconnect Express (PCIe) card was born (see Figure 2-4). In less than ten years, this technology has become one of the most successful solid state storage devices in the marketplace.

Most servers now include PCIe high-speed connections as part of their internal architectures. Integrated circuit boards or cards of various physical sizes can connect directly into the main server circuitry or bus through slots with certain numbers of connection pins. At first, PCIe cards were only that — boards or cards with PCIe connections that held large onboard flash chip arrays. PCIe cards install directly into the server enclosures, so they eliminate SAN network latency. Systems administrators load software that works with the operating system to manage the PCIe cards and create substantial pools of flash-based storage that can be used to accelerate the performance of latency-sensitive applications.



Figure 2-4: A typical PCIe card.

Multiple PCIe cards can be installed in a server, depending on how many slots are available. Over time, a wide range of software and alternative hardware configurations have been developed to address many different IT infrastructure requirements and challenges. They all provide lower latency than traditional SSDs, especially SSDs deployed in SANs, and yet they also offer a similar advantage enjoyed by SSDs — comparatively lower unit prices.

Though PCIe configurations offer almost all the options of other solid state storage devices, most often they're deployed to create large pools of in-server storage that isn't quite as fast as DRAM but much faster than disk-based storage, no matter where the disks are physically located. The pool or cache of very fast storage is managed by software that monitors the activity of the data sets used by the application(s) hosted on the server. The most active data, which will be the data that can most benefit from ultra-low latency, is copied to the PCIe card, and the application then reads it from there, not from other, slower storage.

The PCIe model of data storage is sometimes described as *server-centric application acceleration*. It can offer a range of benefits. These flash products cost much less than today's DRAM while offering memory-like performance. Because the space inside a server enclosure limits their size, you can save on purchase or capital expenses over some other solid state storage products, and because they're made from flash chips, you save operational costs. They can be targeted at accelerating a single mission-critical application. And certainly PCIe cards solve the latency challenge better than any other solid state storage devices.

Their sales in the marketplace have skyrocketed over the past seven years. Nonetheless, both PCIe-based storage and SSDs still compete with the original type of solid state storage device — the standalone appliance.

Solid State Arrays

Before SSDs were invented and long before any PCIe card existed, there were standalone solid state storage appliances. Contrary to some predictions of a few years ago, deployments of flash-based solid state arrays have been substantially

increasing, rather than declining due to the competition from PCIe cards and SSDs. Figure 2-5 shows you one of the current solid state storage array products available in the marketplace.



Figure 2-5: Example of a current solid state storage array.

Although solid state arrays (SSA) were once entirely composed of RAM, now those are rare and built for special-use cases. Instead, the vast majority of SSAs these days are flash-based, and they come in a wide variety of shapes and sizes. At the highest level, the SSA family tree splits into two major branches — those products with close ancestral ties to the operating systems, software, and hardware architectures of commodity servers and conventional hard disk arrays, and the separately engineered, free-standing enclosures or appliances.

For example, you can take a commodity server, load it with an operating system and other software designed or optimized for flash, then fill its disk bays with SSDs — and you have a solid state storage appliance. You can achieve essentially the same effect by mating an enterprise storage array's controller with a disk enclosure filled with SSDs. These are both members of the former branch of the SSA family tree.

The latter branch — purpose engineered boxes of flash chips — are products that use less commodity hardware and modified software, though many still use some. Most often, a custom chassis is filled with units of flash storage — either SSDs, PCIe cards, a memory module called a DIMM (dual in-line memory module), or custom-designed modules. There are versions of this design that also include hard disk drives to lower costs and boost storage capacity.



No matter which branch of the family tree, SSAs aren't designed to be pushed into a server's disk bay or plugged into the PCIe bus. They stand alone. And that, frankly, is their advantage. This is fast storage that is meant to be shared.

They can be directly attached to a single server via one type of interface or another, but more often SSAs connect to application hosts within a SAN architecture.

Shared storage offers many benefits; shared *fast* storage offers many more. Obviously, SSAs offer a much easier way than SSDs or PCIe cards to provide low-latency, high-performance storage to multiple servers — if you already have a storage area network deployed. Though solutions are available, in general it's much more complex to configure ways to share the performance and storage capacities of in-server SSDs or cards than it is to share storage from an SSA. Enterprises that want to connect many servers, often groups or clusters of such, and dozens or more applications to terabytes (TB) of flash storage can do so much more easily with an SSA than if they had to open every server and plug in PCIe cards or SSDs.

Because standalone flash arrays aren't constrained by the physical enclosure of a server or disk bay, nor by the hard disk-oriented interface protocols used in those cases, they can contain a lot of very fast storage. Until recently, most SSAs were rather dumb boxes, and the deployment options almost always necessitated some connection to management or controller devices. This wasn't altogether a bad thing. It was rather simple to hitch them to your SAN by making them part of a storage array or by leveraging the storage management functionality that has been evolving within host side operating systems and related software.



The liability of having data streams travel through networks and into and out of SANs may be overblown. SANs made with Fibre Channel networks add only a few microseconds of latency to the data's round trips. In many cases, the design of the software application itself, or any of a number of other hardware and software components lying in the data path, can add much more latency than the SAN. This fact probably accounts for much of why SSA sales are so rapidly accelerating; even if connected in some networking fashion, they still offer extremely low latency, and the simplicity of sharing their resources is very attractive.

Plus, when you add storage through the network, solving data protection and disaster recovery challenges becomes much simpler. With your data storage in a separate pool, you can easily copy it, then send the copies off to another machine,

another building, or another city. When a hurricane sweeps in off the Gulf and swamps your data center, your mission-critical data isn't lost. And finally, the SSA boxes are becoming a lot smarter now days. This is the big new trend in all of enterprise data storage — flash with all the bells and whistles.

Chapter 3

Choosing Flash Storage Arrays

In This Chapter

- ▶ Analyzing your storage needs
- ▶ Performing tasks with flash controllers
- ▶ Realizing the issues with SSDs and PCIe cards
- ▶ Establishing the benefits of flash arrays
- ▶ Introducing IBM FlashSystem arrays

The decision to actually purchase and deploy flash storage to support your enterprise has two parts. First, you must assess your actual need. Then you must find the solution that fits best. In this chapter, you first look at some of the storage system analytic tools and resources available to help you make the most accurate assessment possible of your storage needs. Then you evaluate each of the flash-based storage options — SSDs, PCIe cards, and flash arrays — and find out why flash arrays are strong candidates to address many enterprise data storage requirements.

Storage Analysis

After you see the need for quicker, more in-depth decision making, faster customer service, or a more palatable data center budget, your next step is to accurately analyze your IT infrastructure to identify exactly what kinds of system performance issues you're experiencing and where specifically they lie. A key to lowering the risks and increasing the value of your flash storage deployments is to thoroughly understand

your system and application performance characteristics to pin-point where and how flash can offer the greatest value.

Most operating systems (OS) offer system monitoring and diagnostic software programs or tools. Two of the most well-known over the years have been Performance Monitor (perfmon) for Windows and `lostat` for the Unix family.



If you have a Unix-flavored operating system such as Linux, use the utility `lostat` to perform analyses of your storage system's performance. `lostat` is a computer system monitoring tool within the Unix family used to collect and show operating system storage input and output statistics. It is often used to identify performance issues with storage devices, including local disks or remote disks accessed over a network.

Specific applications, such as databases, also offer tools that can help you better understand how your computer systems, especially the storage devices, are operating and if there are places within your hardware or software that are creating problems or what are often called “bottlenecks” in the performance of these systems. On the database side, the most well-known of these monitoring and diagnostic tools is Statspack within the Oracle Database application. It is now called Automatic Workload Repository (AWR). Reports generated by Oracle AWR provide database administrators (DBAs) with detailed information concerning a snapshot of database execution time. This snapshot furnishes statistics on wait events, storage input and output volumes, and timings, as well as various views of memory and activities associated with software instructions to the database called SQL.

The statistics and insights provided by tools such as Oracle AWR reports, as well as `lostat` and `perfmon`, about the memory, input and output (I/O), and SQL performance characteristics are invaluable aids in determining if databases or other applications and systems are functioning optimally. From this type of information, you can make much more informed decisions about if your IT environment can benefit from adding flash storage, where specifically your performance issues lie, and even some strong hints about what kind of flash storage product might provide the greatest value.

But frankly, though your own systems administration staff may be very familiar with and use these types of monitoring and diagnostic tools, when you begin to seriously consider

deploying flash storage, you can take advantage of as much help and expertise as you want — furnished by the flash storage product vendors themselves.

All the legitimate flash storage solution providers in the marketplace maintain technical experts often known as Sales Engineers (SE) whose job it is to help you move successfully along this path of information gathering, analysis, solution design, testing, and deployment. Many of the larger product vendors have made significant investments in these types of resources. For example, just in the past few years IBM has invested in laboratories around the world called Flash Centers of Competency (CoC) where potential customers can get in-depth assessments that eliminate risks while maximizing the deployment benefits of flash.

IBM Flash CoC teams offer comprehensive services to potential clients that involve detailed system and application workload analyses called Data Pattern Assessments. IBM experts utilize data center analytics tools to execute end-to-end array, host, database, and file scans of customer environments. IBM Data Pattern Assessments require minimal investments of time and resources from customers, but they return a rich trove of information, which can greatly help you determine valuable data points such as which applications, servers, and storage volumes within your specific IT environment are most impacted by storage performance bottlenecks and unacceptable latency and exactly how connecting your applications to flash storage can provide the greatest benefit.

Flash Controllers

In order for flash to be viable as an enterprise-grade storage medium, several peculiar personality quirks of flash must be mitigated and managed. This is the job of small processors embedded within every flash storage product — *the flash controller*.

Flash controllers perform many tasks associated with writing and reading data to the medium and managing various engineering solutions that help make the particular flash product faster, more reliable, and much longer lasting. Two of the

most common flash management tasks are known as *wear leveling* and *garbage collection*:

✓ **Wear leveling** in enterprise flash storage devices essentially is the activity of spreading data evenly among flash cells to increase flash life. A tremendous amount of very innovative engineering has been focused over the past decade or so on flash controller technologies in order to optimize the useful life span or endurance of flash chips. Unlike consumer uses for flash storage, such as in smartphones or digital cameras, enterprise use cases for flash are characterized by a high number of program (write) and erase (P/E) cycles. Flash would never work in enterprise environments if any particular cell was hit repeatedly with new erases and writes. It would wear out much too quickly. So enterprise flash product vendors design unique wear leveling solutions into their flash controllers to spread the P/E activity out over the hundreds of thousands of flash cells in each device. Wear leveling has become so effective that now enterprise flash storage wears out less frequently than mechanical hard disk drives.

In a recent economic value validation performed on IBM FlashSystem storage, the industry analyst ESG estimated that mechanical disk drives wear out at a rate of approximately 5 percent over a three-year period, whereas flash modules in an equivalent IBM FlashSystem array wear out at a 0.1 percent rate.

✓ **Garbage collection** deals with the performance bottleneck that occurs due to the need for a flash cell to be erased before it can be written to. To cut down on flash cell write times and make flash as fast as possible, flash controllers remember where invalid data exists, such as data that has been updated elsewhere or deleted. Then in the background, the flash controllers erase the cells that contain invalid data and make them available for the next writes coming into the device.



Wear leveling and garbage collection, among many other flash management tasks, aren't accomplished exactly the same or equally well from storage vendor to storage vendor. The speed, latency, consistency, predictability, reliability, and efficiency metrics of each product provide some indication of how well the flash controllers inside it perform. When you load your mission-critical data into the flash storage in which you've just made a significant investment, these attributes will grow very important to you.

Looking at the Challenges of SSDs

After you accurately assess your data storage requirements — how much storage capacity you need, how fast your data must go to meet your business needs, where you have performance bottlenecks within your system, and many other questions — you can begin the process of evaluating your flash storage options. I've already mentioned some advantages and benefits of SSDs, but in this section, you look at some of their challenges.



Solid-state drives have been viewed as the most convenient and lowest-cost way to get flash into your system. But SSDs do have their limitations and liabilities:

- ✓ SSD form factors are very limiting. For example, wear leveling across 10TBs works better than wear leveling across 1TB because you have more flash cells between which to spread out the writes. With a “box of flash” such as a flash array, you can simply make the box bigger if you need more flash capacity to achieve your objectives. With a product designed to fit into the drive bay of a server, this isn't possible.
- ✓ Because they are, in fact, intended to be deployed into the same spaces as hard disk drives, they must use the same interface protocols — essentially networking languages — as hard disks, and these protocols usually were not designed for the speed of flash. In general, all the technologies built around disk drives work well with latencies in the millisecond spectrum, from a few to hundreds. Flash, on the other hand, operates in the microsecond spectrum, ten to a thousand times faster. Some components, either hardware or software, built originally for disk speeds just can't go at flash speeds. So when you deploy flash in those environments, at least some of the potential benefit of flash is wasted.
- ✓ SSDs are rarely the whole solution by themselves. By their nature, they're intended to be plugged into some larger device, whether a server or a storage array. This larger system may not be optimized for flash, and so you pay good money for bad performance.
- ✓ The original advantage of SSDs was purchase price. That's still their advantage, but the enterprise market is growing ever more knowledgeable about costs, and the cost of storage involves more than the purchase price.



Operational expenses count too, such as the cost of electricity, the expense of cooling all those electronic devices, the costs of software to manage the SSDs and the devices that contain and manage the SSDs, and even the cost of data center floor space. Other costs crop up. If you push a bunch of SSDs into your servers, what are the costs of implementing a solution to copy and protect the data on all those individual SSDs? Because they don't do wear leveling as well, how much sooner will you be replacing them?

With SSDs you make cost trade-offs — lower purchase price versus lower overall costs — and you make performance trade-offs as well. SSDs aren't optimized for performance and because they only operate for you as part of something else, a server disk bay or a SAN array brings extra baggage in the forms of complex data paths and added software that degrade the latency and overall performance of the resulting storage solution.

- ✓ Then there is the cost that everyone struggles to define — the value of performance. SSDs, as a type of storage device, are optimized for purchase price and convenience, otherwise they wouldn't use that shape and those protocols. If you want the most performance per dollar, you don't deploy SSDs. In fact, if you want the lowest price per TB of capacity, you don't buy SSDs. As every cowboy knows, ponies are cheaper than thoroughbreds, but you don't take your pony to the horse race.

Understanding the Liabilities of PCIe Cards

Just like SSDs, PCIe cards live by and at the same time suffer from what in fact they claim to be — server-centric application acceleration. The concept hails from the days when one application was hosted on one server. To make that one application perform better, put faster storage in that one server. Enterprises flocked to the idea.

But what about if you used multiple servers to host this application? Server clusters, as these groups of computers are sometimes called, became the first big engineering challenge for PCIe cards. Yet at the same time, server virtualization was gaining traction. This computer architecture involves loading

certain software onto the server to give it multiple personalities. It pretends to be many computers instead of one, each with its own OS and each able to host its own applications. At first, PCIe cards thrived in virtualized environments, until those environments included server clusters.

Of course, engineers soon provided various solutions to the basic issue of sharing the data cached on PCIe cards segregated in their individual machines. But one way or another, the solutions involved networks, and now you had . . . networked storage, which was exactly what the original concept intended to avoid. Plus, implementing and managing all this sharing between PCIe cards in different physical machines involved lots of software, which to the extent that it intruded into data paths and distracted CPUs thwarted the original ultra-low latency objective.

Add to these complications the added labor of needing to pry open each individual server enclosure to install and maintain or replace each PCIe card. Then mix in the limited physical space available within many of these server enclosures, restricting the size, capacity, and capability of individual cards, and enterprises have begun to seriously ponder PCIe cards' ratio of value to complexity.

Establishing the Advantages of Flash Arrays

At this point, you can see that the directions where IT is headed aren't necessarily advantageous for SSDs or PCIe cards. SSA sales themselves and their market share of overall flash sales tend to bear this out. But there are two larger trends that will swamp these smaller differences between storage devices.

- ✓ First, everything related to IT is growing, expanding, accelerating ever more rapidly. More data volume, higher data velocities, more applications, more types of workloads such as mobile and social systems of engagement, more of everything IT.
- ✓ Secondly, software is growing smarter and this is enabling the virtualization of all IT components. Essentially, this is what Cloud computing means — CPUs are becoming

a resource, networking is a resource, and yes, storage is also becoming a resource to be managed, consumed, migrated, updated, scaled up and scaled out, and even subscribed to — all separately from the other components.

How do you transform storage into such a resource with the least amount of headache, complexity, and cost while at the same time optimizing what you want from it — capacity and performance? Enter the flash storage array.

Why you would choose to deploy flash storage arrays instead of SSDs or PCIe cards includes many reasons:

- ✔ **Cost:** The number one topic on everyone's mind is cost. SSDs offer the lowest purchase price; flash arrays offer the lowest total cost, when both purchase price and operational expenses are considered. Plus, flash arrays offer the lowest price per capacity, if for no other reason than there is less packaging per TB. You must buy many separate SSDs to get 50TB of flash storage. You buy one flash array a little bigger than a pizza box. Plus, you must plug all those SSDs into something — something bigger and more expensive than them, and almost certainly not faster.
- ✔ **No slaves to multiple masters:** If you're going to implement networked storage, why not implement the least costly and complex solution possible? Flash arrays are simply optimized to share capacity and all their other attributes with any and all applications that interface with them. If you add more servers, fine. If you add more storage, or reconfigure it, who cares? Flash arrays don't.
- ✔ **Latency:** If you want to accelerate the applications hosted on one physical server, PCIe cards offer lower purchase prices than flash arrays and better performance than SSDs. If you already have or plan to implement shared storage, then you can't beat flash arrays.

IBM FlashSystem

Before you make your decision about which flash storage solution to deploy, I have one more bone to pick. In the previous chapter, flash arrays were introduced as descending from a

primal ancestor along two family tree branches — flash arrays with many components, especially their hardware, not purpose-built explicitly for the role of flash array, and flash arrays with all their components, essentially, purpose-engineered for this one role.

Why the difference? Cost, of course — the cost of development and the cost of deployment. The truth is, most of the flash arrays on the market today began as software engineering projects. A few ingenious software engineers built new software that solved a particular storage problem. Then they loaded it on hardware they essentially bought at the store. If everything worked as planned, they had a competitive new flash array product to sell, the result of relatively low-cost development coupled with quick time to market. Good business. But, good for you? Not necessarily.



If you want to combine the cost of flash and disk in your storage deployment evaluation, simply buy the best of both and integrate them with readily available automated storage tiering software. You probably already have a disk-based storage array; just add more of the lowest cost disk you can find. Then deploy an IBM FlashSystem array. This solution offers lower cost per TB and much higher performance. It also provides excellent storage virtualization and management software, including dynamic tiering that automatically moves data between storage media based on the policies set by you. Figure 3-1 is a photo of an IBM FlashSystem V9000 with all the storage bells and whistles included.



Figure 3-1: IBM FlashSystem V9000.

With the disks in their favorite environment and the flash highly optimized within its IBM FlashSystem chassis, you can add, change, and evolve this storage solution to meet your evolving application and business needs without affecting any other component of your IT infrastructure, including other storage devices, and without ever throwing any disk or array component away until they fail.

IBM FlashSystem all-flash storage technology is purpose-engineered from the circuit to the chassis for the future of information technology and of business itself. The hardware descends from solid state storage ancestors going back literally decades. This is eons in IT time. The software in the base IBM FlashSystem models can make the same claim. For the model that replaces traditional enterprise storage arrays, IBM has tightly integrated IBM's industry-leading storage services and virtualization software into the IBM FlashSystem mix. This software comes from a suite that has been deployed successfully in thousands of demanding IT environments over the past decade. For example, the IBM Real-time Compression function alone is based on over 70 patents. Then, IBM research and development labs around the world are continually working to improve and enhance both IBM FlashSystem hardware and software.

If you want to deploy flash storage at the lowest \$/TB, look long and hard at IBM FlashSystem. If you measure cost by \$/performance, then beating IBM FlashSystem won't be easy. Lowest possible latency? Check. Need to start small and then grow your flash investment as your budget grows? Check. Don't have the manpower and need a solution that's especially easy to deploy? Check.

Most importantly, IBM FlashSystem means that you can finally move on from disk, and it's cost effective. You can stop writing applications and architecting compute environments to mitigate the shortcomings of traditional storage. You can truly tackle the problem of endlessly rising power consumption. You can fully exploit the potential of Cloud computing and big data or implement virtual desktops without your storage infrastructure getting in the way. In fact, at the end of the day, that's really the biggest IBM FlashSystem benefit: Its performance, reliability, and efficiency turn storage from a limiting factor into a real driver of innovation within your business.

Chapter 4

Exploring Deployment Designs with IBM FlashSystem

In This Chapter

- ▶ Learning how to directly attach IBM FlashSystem
 - ▶ Using SANs to your advantage
-

You are an IT decision maker in a vibrant enterprise, perhaps an e-commerce business, a hospital complex, or an academic institution. You've familiarized yourself about data storage technology and that journey led you to IBM FlashSystem. Now, you must design and implement a successful deployment strategy. Of course your business needs, and how your current IT environment addresses and supports them, guide you. You've already tapped into a worldwide network of Flash Centers of Competency, Lab Services, Sales Engineering teams, and solution architects who've analyzed your needs, helped formulate the best solutions, and provided resources and guidance for both off-site and on-site proof-of-concept testing. To begin the process of developing the most effective solution architecture, you and your IBM team have to evaluate your various deployment options. This chapter introduces the basic flash array deployment architectures and provides some thoughts about why you may choose one over the other.

Deploying Direct Attached Storage with IBM FlashSystem

If you think that SANs are the only way that enterprises deploy IBM FlashSystem arrays, think again. What if you're one of the hundreds of businesses around the globe that directly engage in or provide IT support to online equities trading? One application essentially defines your business. Milliseconds define your timeframes. Moving at literally lightning speeds and yet being able to guarantee the legally required capture and ultra-reliable storage of every transaction are business requisites. Or you process data received from a weather or research satellite? You analyze old seismograph data looking for hidden petroleum reserves. You're installing smart meters for five million utility customers, and this one application processes the multiple data streams from each meter to better manage your portion of the power grid. Direct attached flash storage may mean the difference between discovering a new planet, increasing your revenue by millions, or a blackout.

The *direct attached storage* (DAS) architecture is a venerable, fairly simple storage solution design still used in many IT environments, not only for business reasons but also for technical reasons such as when connecting storage to large computers called mainframes. Essentially, DAS refers to storage architectures where the storage device is linked directly to the application host(s) with no or minimal intervening networking resources. Normally, when you deploy a DAS storage solution, you don't include a network switch, but instead cable the storage device directly to the application server. You can see an example of what DAS looks like in Figure 4-1.

Directly attaching IBM FlashSystem arrays is pretty simple. Just follow a few steps:

1. Install Host Bus Adapters (HBA) in each server where you want to directly attach IBM FlashSystem.

HBAs are hardware components with some software that enable servers to interface with networks or send signals directly through appropriate cables to another device.

2. **Connect the HBAs to the IBM FlashSystem ports (ensuring failover across all potential components) perhaps with Fibre Channel cables or Ethernet and carve capacity from the flash to present to the server.**



If the array is serving a single application or server cluster, you can utilize the Open Access model of IBM FlashSystem with enhanced management features to simplify creating logical storage volumes called LUNs. Access to IBM FlashSystem LUN provisioning and configuration is automatically open to all connected servers.

3. **Install the appropriate multipath configuration on each connected server's OS and utilize the new volumes as if they were any traditional disk.**




Figure 4-1: The simple DAS architecture.



Directly attaching IBM FlashSystem storage has many benefits:

- ✓ **Lowest possible latency:** HBAs directly cabled to other appropriate devices generally add only 10 to 20 microseconds (ms), which is only 10 to 15 percent of the latency of the IBM FlashSystem unit itself.
- ✓ **Greater control:** This is dedicated storage; you don't have to share it with anyone. Its performance and capacity are all yours if you own the application.

- ✓ **Higher capacity:** You can enjoy memory-fast storage speeds without the physical capacity limitations of in-server solid state storage.
- ✓ **Greater reliability:** Standalone arrays like IBM FlashSystem offer better data protection, resiliency, and serviceability than in-server cards or SSDs.
- ✓ **Better failover:** Server clusters and the software that manages this configuration offer the crucial “fail safe” advantage of supporting “hot” backup virtual servers or virtual machines (VM). If a physical server fails, all its VMs can instantly migrate to their backups on other physical servers. Your business critical applications should never know that a server failed. But if a physical server fails, so do all of its in-server storage devices. Direct attached IBM FlashSystem has Active/Active ports and controllers for each LUN, so as your VMs fail over, storage remains available and unaffected — just like your business.
- ✓ **Easier upgrades:** With storage separated from servers but not configured behind a Storage Area Network (SAN) switch, you can add, subtract, and upgrade any IT component without affecting or needing to upgrade others. For example, you could upgrade your Fibre Channel cabling to increase the bandwidth to your storage and you wouldn’t need to install new networking switches or other IT infrastructure.
- ✓ **Greater data security:** Less infrastructure complexity means less risk of failures that can affect your business.
- ✓ **Lower infrastructure costs:** Just like with security, less infrastructure complexity leads to lower costs.



It’s also important to point out the liabilities of DAS architectures:

- ✓ **Less flexibility:** Depending on how all your applications are hosted, it may prove more difficult to share direct-attached IBM FlashSystem storage with all of them.
- ✓ **Lower utilization:** It’s quite possible that IBM FlashSystem performance exceeds that of the servers to which you’ve connected it, leaving some of its capabilities untapped.

Connecting Systems with a SAN-Attached Architecture

The vast majority of enterprise IT infrastructures these days are moving toward network attached storage, if they already hadn't years previously. Why? For many good reasons, you don't want a server; you want application hosting/data processing resources where, when, and how you want them. The same holds true for enterprise storage. Applications and their owners and users don't need to know about or care how you do it; they just want to get at their data when, where, how, and how fast they need it. To enable this revolutionary functionality, storage must be connected to the other IT components through networking.



A powerful aspect of SAN-attached architectures is that you can connect most any kind of appropriate storage system. You will have your IBM FlashSystem array of course. You can also connect hard disk drive arrays, often called *RAID systems*, and even devices that write data to CDs or tape, usually for backup or archive purposes.

To deploy IBM FlashSystem storage, you just install the array(s) in a rack in your data center and connect it with the proper cables (again Fibre Channel is the most popular) to the SAN's management software, often referred to as the *name server*. The management software in the IBM FlashSystem array introduces itself and makes the interfaces or ports visible. Most often, you then implement a feature called *zoning* in the Active/Active mode, and IBM FlashSystem storage, performance, and features will be available to your applications.



Your IBM FlashSystem storage becomes a resource, available to any device connected to the opposite side of that SAN switch. But there are many other advantages of the SAN deployment architecture:

- **More options:** SAN deployment is the most common model, which means there are more software and hardware products available to support and extend it. For example, almost all major switch brands and models are qualified to support flash storage deployed in a SAN.

- ✓ **Easier expansion:** A SAN enables more flexibility to expand storage capacity and/or performance because the same server ports can connect to multiple storage systems.
- ✓ **Greater share-ability:** Any application/VM/server connected to the SAN can leverage a portion of the IBM FlashSystem resources.
- ✓ **Larger clusters and data sets:** More and/or larger server clusters can be implemented and data sets of all sizes can be shared between multiple applications, server clusters, or replicated around the world.
- ✓ **Easier scaling:** There's essentially unlimited room to grow and tailor storage resources to match application needs.
- ✓ **Higher throughput:** Server clusters can use the full throughput capabilities of the SAN when needed for failover scenarios or spikes in application data traffic.



SAN-attached flash storage does have a couple of disadvantages:

- ✓ **Slightly lower performance:** Because of the networking involved, some, though minimal, performance impacts are inevitable relative to direct-attached and in-server solutions. Often this cost in latency is easily offset by the greater flexibility of capacity/performance scaling and easier/simpler deployment and maintenance.
- ✓ **Higher switch performance requirements:** When you deploy fast storage, all the other components in the data path must be optimized to handle the increased performance levels or your investment in flash will have less impact.

Chapter 5

Implementing the Future with Virtualized Storage

In This Chapter

- Explaining the advantages of storage virtualization
 - Virtualizing all your storage using IBM FlashSystem
-

Virtualization is the future of enterprise data storage. You focus on your business; your storage virtualization engine focuses on increasing the efficiency, performance, security, and accessibility of your data, while lowering its cost. Storage virtualization offers some serious insulation against the future. New storage media technologies have already galloped over the far horizon of possibility; they just haven't climbed the nearer hills of cost-effectiveness yet. But if your storage is virtualized, and you view, manage, and consume storage as a resource, no matter what the particular storage medium happens to be, your virtualization engine will manage it appropriately while your applications, and most importantly your business, will never know the details.

IBM FlashSystem offers storage virtualization deeply integrated with ultra-fast flash — dozens of TBs of capacity all in an enclosure the size of a few pizza boxes managed with ever-increasing intelligence and sophistication, traveling at the speed of integrated circuitry. You can purchase IBM FlashSystem models that include the technology necessary to virtualize all or any parts of the storage systems that comprise your SAN. After you deploy the array(s) and its ports are presented to the name server, look to the easy-to-use IBM FlashSystem graphical user interface (GUI) to virtualize your existing storage under one management pane of glass, so to speak. With this capability, you can then extend the powerful

suite of IBM storage management features to all the other storage systems in your SAN.

Virtualizing your storage using IBM FlashSystem changes your SAN from a group of storage systems into an IT infrastructure resource that can be allocated, reallocated, scaled up or down, upgraded, and on and on with no impact to and not even any awareness of such by your many applications.

Storage virtualization offers benefits in a number of areas at the heart of enterprise storage, including pooling, tiering, data protection, and data reduction/capacity optimization:

- ✓ **Pooling** brings storage resources together so that the appropriate capacity can be delivered to each application, and the magic of reallocating these resources is enabled.
- ✓ **Tiering** brings storage resources together so that the appropriate performance can be delivered to each application.
- ✓ **Data protection** involves the multiple ways that enterprises ensure against data loss or corruption.
- ✓ **Capacity optimization** most often utilizes “thin provisioning” plus various data capacity reduction technologies to reduce the amount of idle or redundant data stored and managed by your storage system, saving you money in several ways.

Storage Pooling and Tiering

Essentially, storage virtualization enables automatic matching of application workloads to the right storage resource. Before storage virtualization, the data used by a particular application, called its data set, was stored on a specific physical collection of hard disk drives. To move that data set to another storage resource, the application had to be turned off, then all the information was moved, or migrated, the application was updated with the new physical locations for its data, and then everything was turned back on. The same might happen if you wanted simply to add more storage capacity, say because your application was growing — shut things down, reconfigure, spin back up. It was very expensive, in terms of time and

labor, but also because of the business or operations productivity lost while the application was offline.

Virtualization based on disk storage has enabled automatic data migrations. But because fetching data from disks and writing it to other disks is slow (certainly compared to other IT components), these data migrations are more like the movements of elephants or buffalo across the Serengeti Plain. Virtualization enables them to happen without direct impact to applications, but they don't happen fast.

Hitching flash to storage virtualization changes things dramatically. Flash transforms data migration into data mobility. Leading edge flash-based storage virtualization such as that integrated into IBM FlashSystem can move entire data sets but also just portions of data sets, volumes, and sub-volumes from one storage resource to another very quickly. Now it's the gazelles darting and leaping, instead of the elephants plodding.

For example, some data sets become active only at certain times — think of month-end accounting applications. And perhaps only certain portions of them. With virtualized flash, to use the phrase that's becoming fashionable, the virtualization engine is constantly monitoring data activity and when the end of the month rolls around and parts of that data set become active, they can be moved quickly, automatically, transparently from storage optimized for capacity, such as disk or tape, to storage optimized for speed — flash.



A SAN can include multiple storage media. This is the way SANs normally evolve. For discussion purposes, I'm going to assume you started with a SAN composed of a single disk storage array. Over time, you added another, or others, as your business grew and/or diversified. Finally, performance and cost factors caused you to add a flash array to the mix. If it was the appropriate IBM FlashSystem model, then you could bring all your separate storage systems together so that to your applications they appear as a single pool of storage resource. Within this pool of storage resource, your disk systems will be slow but relatively inexpensive per unit of bulk capacity, and tape-based systems will be even more so. Your flash will be less expensive per unit of performance, much less. And now you see the rationale for what is called tiering. To lower costs and increase performance and efficiency, you place data on the most appropriate storage medium.

Virtualized storage does this automatically, continually seeking to maximize the utilization of your various storage resources, whatever they may be, based on whatever policies you set as its priorities. Virtualized flash does this with greater agility, leading to even greater savings and efficiency.

Data Protection

Another crucial function that your storage solution must perform in some manner is *data protection*. Essentially this means that when your applications request it, your data is available, and if some component or process fails within your IT infrastructure, none of your data is lost forever.

Data protection actually becomes a very expensive proposition and is usually approached in two ways — preventing failures from happening, or at least from affecting the integrity of your data, and making copies of the data so that the copy can be used if the original is lost or corrupted.

To address the former and prevent failures, most enterprises operate by the simple rule — no single point of failure. This means that within the data pathways themselves, if any one component fails, data will not be lost or corrupted. Because hard experience has taught us that nothing is perfect, the only way to ensure that no failure will result in lost data is to make everything redundant. But a minimum of two of everything drives up costs dramatically.



IBM FlashSystem helps you lower data protection costs by engineering the arrays themselves with no internal single point of failure. So, you make multiple redundant connections from the SAN switch to the machine itself and data travels through redundant pathways from the interfaces into the separate, redundant flash storage modules. Even with this level of reliability built in, some enterprises will still configure their storage architectures with whole redundant systems. This type of configuration is often referred to as *mirroring*, or deploying a “hot spare” that can take over if the active system fails. But IBM FlashSystem’s no single point of failure does offer the option to forego the need to mirror or configure spares, and this can result in much lower equipment purchase expenses. As a matter of fact, achieving the no-single-point-of-failure internal array architecture required years of engineering to

accomplish and isn't necessarily available on all other flash storage arrays.

Internal hardware redundancy is not the only way that IBM FlashSystem protects your data. The systems also employ RAID-based data protection regimes. Using this technology, a unit of data is split into several parts and each is written to a separate flash chip within a flash module of the array. Then a key, known as a parity bit, is calculated by the controller and added to the data unit. The parity bit enables reconstruction of the entire data unit, if a flash chip fails and that part of the data is lost.

IBM FlashSystem uses a unique solution called *Variable Stripe RAID* in each individual flash module. This innovation allows the RAID algorithm to evolve if a chip fails, so that other flash resources in the RAID group aren't unnecessarily thrown out with the failed chip, dramatically increasing efficiency and lowering costs.

Then IBM FlashSystem goes another level better; it uses RAID again, only at a system level between all flash modules, instead of just inside each individual module. This means an entire flash module could fail and you wouldn't lose any data. The two data protection components — module-level Variable Stripe RAID and system-level hardware RAID — operate independently, but together they provide synergistic system fault tolerance to mend multiple flash memory failures.

No single point of failure, redundant components and data paths, two dimensions of RAID, and these aren't all the ways IBM FlashSystem protects your data. Individual flash cells aren't all perfect; some hold a charge well and can be accurately read, and some don't. From the beginning of the use of flash in mission-critical environments, flash engineers compensated for the lack of flash perfection with what is known as Error Correction Codes (ECC). ECC algorithms are applied by the flash controllers while data is being read to check for errors and correct them on the fly. IBM uses a proprietary "hard-decision" algorithm to deliver very high correction strength with lower processing overhead. The overall result is that IBM's unique error correction solutions drive up performance, reliability, and throughput while driving down complexity and cost.

Even though IBM FlashSystem provides many layers of data protection within the array itself, in general, storage systems

haven't been and many still aren't so resilient. So enterprises over the years have devised means outside of or not dependent on the particular hardware or system to protect valuable data. The most common of these is to make a copy of it and store that copy somewhere else.

The two most popular ways to copy data sets are called *snapshots* and *clones*. Snapshots involve essentially taking quick pictures of the data set at specified moments in time. Then, if data is corrupted, the system can be moved back in time to the last snapshot and started again with data that was correct at that point. Obviously, the more often you do snapshots, the more recent your backup will be. But of course you must store these snapshots, which takes resources away from your primary application workloads.

Snapshots lead to two challenges that storage virtualization addresses especially well. First, the process used to reduce the storage resources needed for snapshots can cause significant impacts to storage performance because they involve more processing and software in the data path. Storage virtualization can dynamically move snapshot activities out of “the line of fire” so to speak, utilizing storage resources available at any particular moment that will least affect latency. Virtualized flash goes one better — it offers extra performance and lower latency so that snapshots can be done in flash-based resources with minimal impact on overall performance.

Next, think about the situation in most SAN environments — they have multiple systems, some flash, some disk or tape, each different, and often none that “talk” to each other. How can we manually perform a coherent snapshot across all these disparate systems? With storage virtualization we can, because the storage is managed as one resource, not as different systems. Virtualization tools such as IBM FlashCopy Manager can synchronize and manage snapshots across disparate slower and faster arrays and use IBM FlashSystem resources to almost eliminate performance impacts.

Clones are another data protection strategy that storage virtualization enables. Clones are complete copies of the entire data set, very different from space-efficient snapshots that may just capture the changes to data. Clones are used to recover from disasters and major system failures. Another important use is for the software development, testing, and new application qualification environments, which are carefully

segregated from the actual production environments but still need to use a relatively accurate or legitimate version of the data set. Storage virtualization enables data set clones to be “shipped” to separate storage resources whenever needed for software development and testing with no impact on the production environment. Appropriate storage resources can be quickly allocated and configured for these use cases based on the capacity and performance needed and/or available.

Capacity Optimization

Historically, storage capacity has been a static resource. You have this much, period. To add more, you must stop everything, physically haul in and configure more disks or new systems, then spin it all up again and hope nothing explodes. To avoid the risk of running out of storage capacity unexpectedly and to account for growth, you allocate or provision a lot more than you actually need right now. This is called “over-provisioning,” a venerable and expensive storage management practice that can result in a lot of resources spinning happily away unused.

Thin provisioning means allocating only the storage resources you need right now, the opposite of over-provisioning. It’s much more efficient and less expensive, but in traditional storage environments it’s too risky. Not when you’ve implemented storage virtualization. When you can add capacity quickly and easily, the storage world flips.

With virtualized flash such as IBM FlashSystem, when you deploy and configure this new technology and make all of your storage a single resource, you literally try to find ways to allocate 100 percent of the IBM FlashSystem capacity. You want all of that extraordinary performance working for you, right now. IBM FlashSystem comes with thin provisioning technology that, in fact, allows you to over-allocate its capacity. Thin provisioning functionality carefully monitors actual storage usage and automatically allocates more from other LUNs or other available systems just when needed, then allocates it elsewhere when no longer needed. If a call comes in at 2:00 a.m. on Saturday that data volumes are reaching 90 percent utilization and climbing quickly, external storage can be easily zoned in by storage virtualization and utilized for data growth without panic or even a trip to the data center.

Another storage capacity reduction technique is called data compression. Compression is the reduction in size of data in order to save space or network transmission time. Applications write data to storage and during the write process a tool such as IBM Real-time Compression shrinks the amount of storage capacity needed by implementing software- or hardware-based formulas that remove all extra-space characters, insert a single-repeat character to indicate a string of repeated characters, and/or substitute smaller bit strings for frequently occurring characters, among many other techniques. IBM Real-time Compression can reduce certain types of data files by a ratio of up to 5:1.

Because flash is still more expensive per unit of capacity than some disk storage, data compression and other reduction strategies tend to offer even greater benefits when applied to flash storage. When data compression is implemented using software running on commodity processors, it can significantly impact the storage latency. IBM FlashSystem therefore implements IBM Real-time Compression using a mostly hardware-based process, which minimizes the latency impact while maximizing the degree of compression. Also, some data types don't yield much benefit from compression algorithms. During deployment, or at any time afterward, the virtualization engine within IBM FlashSystem allows you to enable IBM Real-time Compression only on the data volumes you specify, thus optimizing their performance.



The flexibility of data compression deployment offered by IBM FlashSystem results from a large group of IBM innovations collectively referred to as IBM FlashCore technology. These various innovations enable IBM FlashSystem to deliver the wide range of operational and cost efficiencies, such as the agility of Real-time Compression. IBM FlashCore technology lies at the heart of FlashSystem storage. Fundamental to this technology is the concept of the hardware-accelerated data stream that delivers very high performance while also supporting the capacity optimization features essential to modern enterprise-class storage.

Because the engineering embodied in IBM FlashCore technology is so strong and yet so flexible, it enables IBM FlashSystem to incorporate new performance and capacity optimization features, as well as many other virtualization capabilities, with absolutely no compromise in system performance or reliability for many years into the future.

Meet your storage challenges with flash arrays

Get the most current thinking about what you should do as the responsible manager or technician if you are assigned the task of implementing a flash storage solution. If you're an IT decision maker, find out why all-flash storage is cost-effective and how easily it can be deployed, configured, and operated.

- *Define data storage-related problems — consider a flash storage solution*
- *Look at various types of flash storage — understand what they're used for and who currently uses them*
- *Get to know flash storage systems — discover the benefits of all-flash storage arrays*



Open the book and find:

- What higher performance data storage technologies are currently available to you
- How to use flash to your advantage
- Why you should choose all-flash storage arrays
- How to deploy IBM FlashSystem
- Why storage virtualization is in your future

Go to Dummies.com

for videos, step-by-step examples,
how-to articles, or to shop!